Title: Development and Testing of Next Generation AWACS SF Processing

Author: Tom Barker, Lockheed Martin, Owego NY
Dr. Larry Holmquist, Lockheed Martin, Owego, NY

Abstract

The development history, system capability and testing of the Lockheed Martin Tracking and Sensor Fusion system is described. The system has been selected for the AWACS platform and is currently in the final stages of flight qualification. The basic algorithm utilized is an extension of the classic MHT algorithm. In addition to the n-dimensional assignment processing used to select optimal solutions for the developed set of hypotheses, the program incorporates a full set of features to address the needs of hard real time scheduling and open system methodologies that facilitate addition of extension processes.

Introduction

Lockheed Martin developers recognized in the early nineties that a significant increase in the computing power available to airborne platforms was eminent. In the same period operators were saying the then state-of-the-art AWACS tracker was inadequate. Given these two factors developers initiated an effort to discover a tracking and sensor fusion algorithm that would fully exploit all available compute power to significantly improve tracking quality.

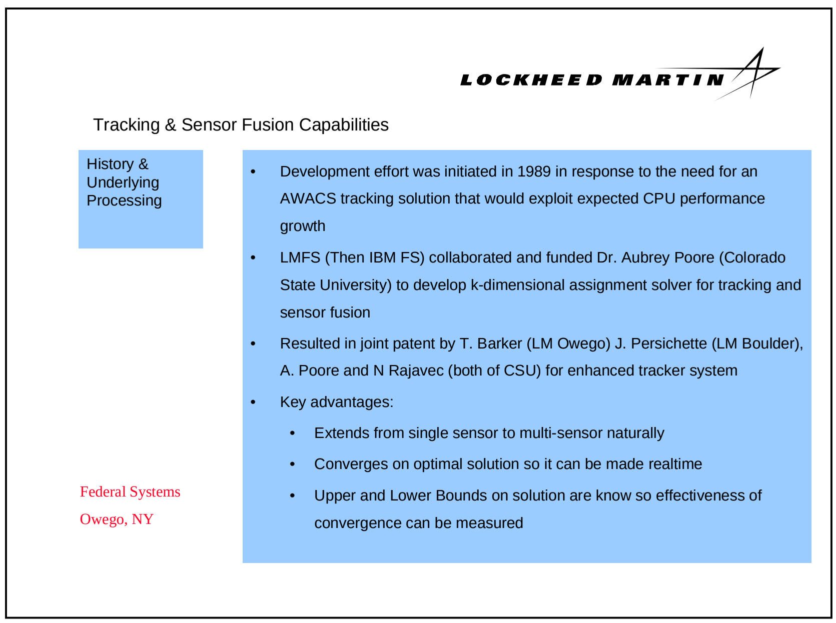

Lockheed Martin developers realized that if a recent development for solving 3-dimensional assignment problems could be generalized to n-dimensions, then it would solve the tracking sensor fusion problem. Although the true n-dimensional assignment problem is NP Complete, tracking and sensor fusion are stochastic and therefore not actually in the same class. That characteristic made the problem seem approachable. Lockheed Martin supported a joint research collaboration with Colorado State University to address the problem. This effort developed a solution to the problem that converges toward the optimal result. A converging solution makes it possible to allocate fixed amounts of processing power and that permits the processing to be placed on a real-time schedule. What we have found is actual problems converge quickly to near optimal solutions for almost every track.[1] If the solver continues to process then it tends to cycle through solutions that are effectively equivalent, given the level of noise in sensor observations.

---

[1] It is possible to construct pathological cases by considering the processing and specifying observations to thwart the solver. When we have done this the resulting track looks absurd. To date we have not seen this occur in real data so we have not incorporated any special testing to detect the condition.

Form SF298 Citation Data

Report Date *("DD MON YYYY")* 00001999	**Report Type** N/A	**Dates Covered (from... to)** *("DD MON YYYY")*

Title and Subtitle Development and Testing of Next Generation AWACS SF Processing	**Contract or Grant Number**
	Program Element Number
Authors Barker, Tom; Holmquist, Larry	**Project Number**
	Task Number
	Work Unit Number
Performing Organization Name(s) and Address(es) Lockheed Martin Owego, NY	**Performing Organization Number(s)**
Sponsoring/Monitoring Agency Name(s) and Address(es)	**Monitoring Agency Acronym**
	Monitoring Agency Report Number(s)
Distribution/Availability Statement Approved for public release, distribution unlimited	
Supplementary Notes	
Abstract	
Subject Terms	
Document Classification unclassified	**Classification of SF298** unclassified
Classification of Abstract unclassified	**Limitation of Abstract** unlimited
Number of Pages 8	

As the solution converges the processing calculates and utilizes upper and lower bounds on the effectiveness of the current solution. That feature when combined with the iterative nature of the processing permits the solution to be optimized to the particular tracking situation encountered. That has been exploited to provide operators with the capability emphasize particular tracks or sectors. The particular optimizations available to the operators are described in latter sections.

The processing developed has resulted in a tracking and sensor fusion processing patent that is held jointly by Colorado State University and Lockheed Martin. The inventors are Drs. Aubrey Poore and Nenad Rajavec (CSU) and Tom Barker and Joe Persichette (LM).

Processing Organization

The primary processing steps are shown in the adjacent chart.

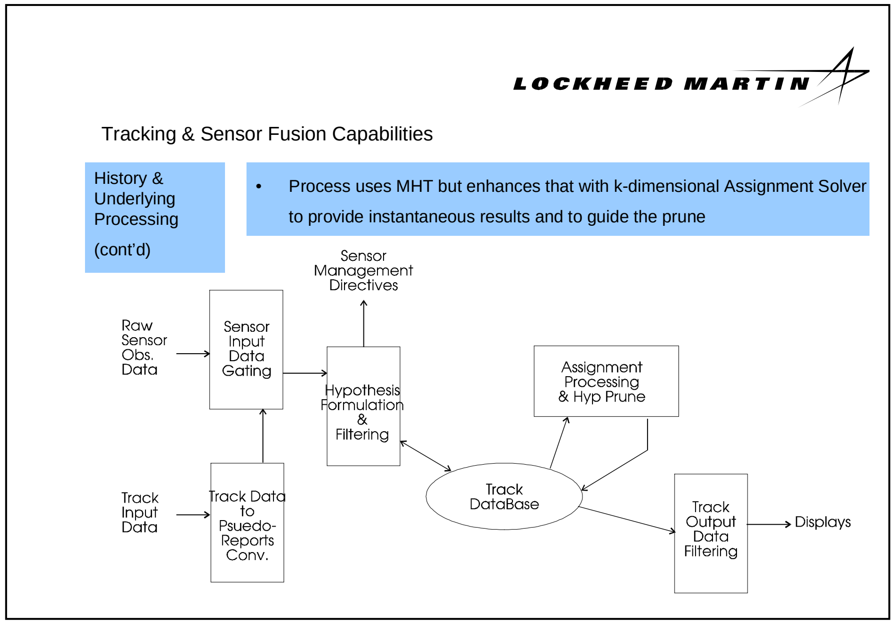

Raw sensor data is initially matched to the track history using an approach in which each observation is matched with any existing tracks that might potentially match. Each resulting hypothesis is filtered and that filtering results in a likelihood metric for each hypothesis.

The output of hypothesis processing is stored in a memory resident database. The database is the source for data passed into assignment processing, output processing and although it's not shown sensor input gating.

Note that there are two types of input to a sensor fusion system. Raw sensor data includes sensor noise that is not correlated over multiple observations. What we have called 'Track Input Data' is a class of data which is correlated. For example, data that is coming from other tracking systems. For this class of data we assume outputs come from a Kalman Filter and as such, observations are a weighed average of track history and current observations. Since Hypothesis Formulation involves Kalman Filtering we include processing to 'de-correlate' this type of data. In effect we use an Inverse Kalman Filter to implicitly generate a pseudo-observation.

Hypothesis Formulation

For illustration consider two tracks, T1 and T2. Based upon their history we can predict the region where each track might be at the next observation time. In a normal tracking situation these regions may overlap and multiple observations might be detected in each region. Given that we create multiple hypotheses to extend the tracks to all potential matching observations. Each new hypothesis can then be projected forward to the time of the next observations.

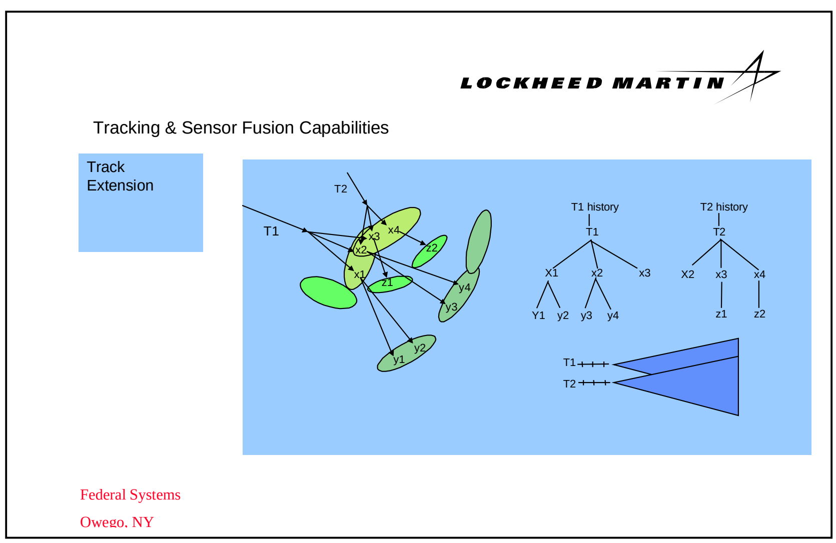

In effect we are building a tree, rooted on the track history, of hypothetical track extensions. Within the tree, levels can represent either consecutive observations from a single sensor or observations from multiple different sensors. In the first case we are dealing with a tracking problem while in the second case the same processing is used for sensor fusion processing.

For graphical convenience the tree structure of a tracks hypothesis set can be illustrated by the simple triangle.

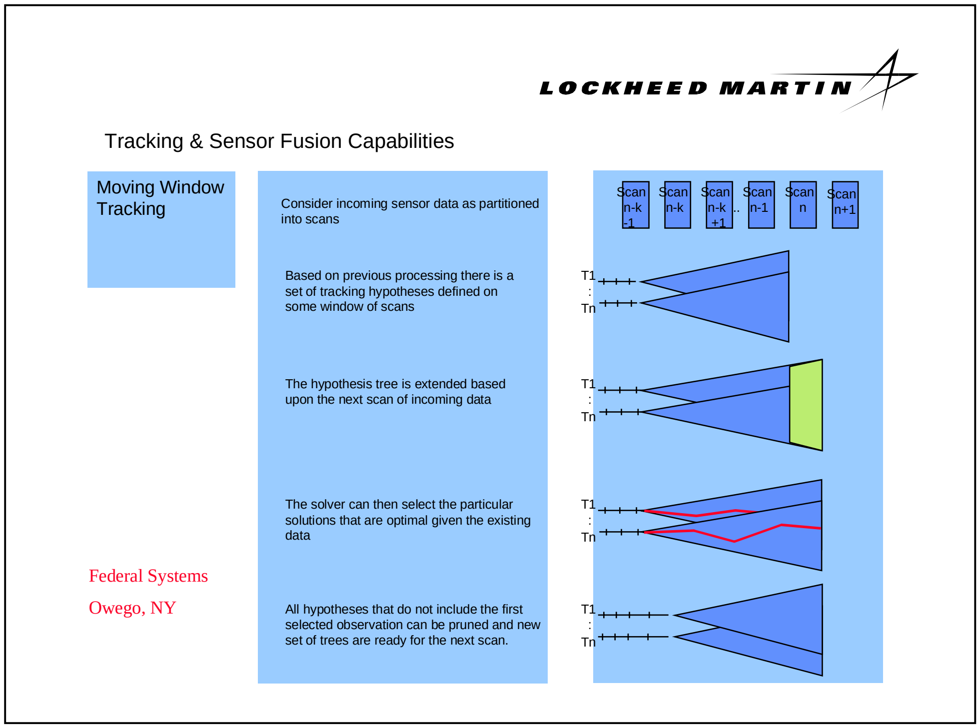

Assignment Processing

The processing can be assumed to start with multiple hypotheses covering a set of n-1 observations sets. In a tracking problem we have found that 3 to 6 observations sets are needed to resolve the ambiguity in observation to track assignments. In the sensor fusion cases we might want to maintain the hypothesis set over perhaps two sets of observations from each sensor. In the tracking case that means we are leading up to a 4 to 7 dimensional assignment problem. Sensor fusion problems frequently result in larger problems. For example, with 6 sensors, and 2 sets of observations per sensor the hypothesis tree would span 12 sets of observations and result in a 13-dimensional assignment problem. (AWACS currently uses 10-dimensional processing.)

As new observations are received the hypothesis tree is extended. The extension need not be for a single set of observations, as is shown in the illustrations. For explanation purposes it is easiest to consider the extension to include all the data from a single sensor scan, but the algorithm does not impose this restriction. In actual processing we extend the trees over a part of all sensor observation sets to minimize latency as perceived by the operator.

Assignment processing selects the subset of hypotheses that serve to assign all observations to a track or to declare the observations to be false or initiating tracks, while simultaneously maximizing the global likelihood of observation to track matching. In the illustration we show the solution as being a particular selected hypothesis for each historic track. The selected hypotheses are (1) the solutions that will be output to the operator and (2) the basis for hypothesis tree pruning. Consider the tree branch that includes any selected observation. It forms a tree rooted at a more recent time than the original hypothesis tree. Pruning is accomplished by simply moving the hypothesis tree froward to the some observation in the selected solution. Depending upon the particular problem this move can be one or more tree levels.

Sensors Supported

The Sensor Fusion capability supported on the AWACS platform includes Radar, IFF, ECM, and crosstold tracks. The prototype solver for Multi-Sensor Integration (MSI) adds ESM and other sensor processing. The Radar and IFF sensor systems are physically located on opposite sides of the dome so reports from the two sensors are 180° out of

phase. ECM reports occur when the AWACS radar is being jammed and they constitute an angle only report on the jammer. In today's environment cross-told tracks are reports being received from other AWACS platforms and fighters. This class of report could easily be expanded to include reports received from intel or other service sources. Both crosstold and ESM reports are treated as reports from other trackers. This data is preprocessed by the Inverse Kalman Filter.

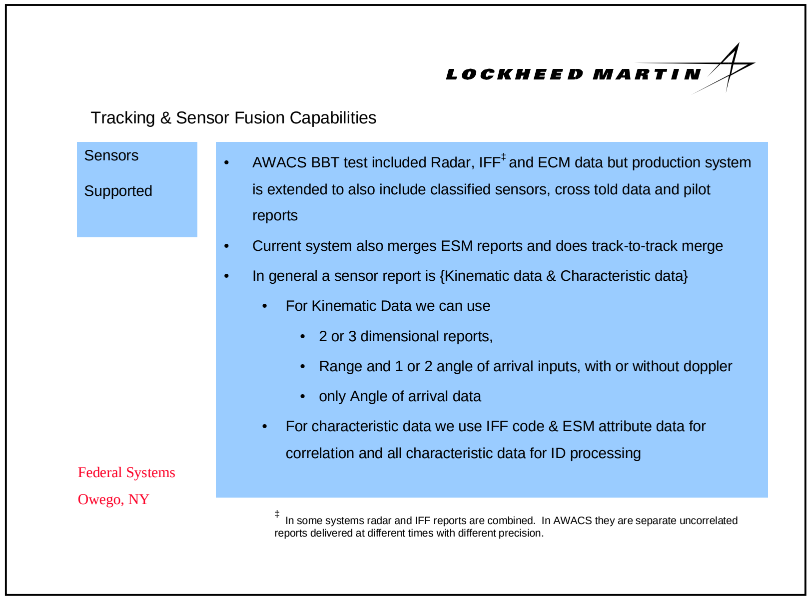

The sensor fusion processing capability is not limited to just the particular sensors used for AWACS. Instead of developing ad hoc processing for specific sensors generic processing for classes of sensor types is supported. This involves two types of processing: kinematic data processing and characteristic data processing.

Kinematic processing uses the location data derived from sensor reports to fit the most likely target motion path to the observed data. This processing uses parameter driven Kalman Filtering that they can be dynamically adjusted to a particular sensor. Parameters include standard deviation for range or azimuth. These can be adjusted as a function of some other characteristic of the report. To address all type sensors the tracking system includes filters for:

- angle only sensor reports
- 2-demensional reports providing range, angle and optionally a doppler component
- 3-demensionally reports that add vertical angle of arrival data to the 2-dimensisonal type report.

These filters share the same state and covariance matrices so they can be interspersed when processing a single track. The filtering process develops a track covariance matrix and from that matrix we derive a likelihood metric for the hypothesis. It is the likelihood metric that will be maximized by assignment processing.

Sensors frequently provide characteristic information in addition to location parameters. For example IFF reports will include the transponder code, ESM reports may indicate a type of transmitter or a signal characteristic like pulse repetition rate. As these reports are combined into a track hypothesis the processing can analyze the consistency of characteristic data within a hypothesis. The results of this analysis modify the likelihood metric derived from kinematic data. For example, consider a series of observations with location data that can fit a flight trajectory. Kalman Filtering will in that situation generate a high hypothesis likelihood metric. But if those observations contain several different IFF codes then characteristic data processing will reduce the hypothesis likelihood.

The number of observations in the hypothesis also impacts track likelihood. This level of analysis permits parameters to be incorporated which relate to global sensor properties, for example, probability of detection and false alarm density.

Output Data

The Tracking and Sensor Fusion SW does not generate operator displays. It functions as a server and will respond to display or other application SW requests for current tracking state. On the AWACS platform the application and display functions request estimated track location and velocity data. To satisfy these requests the tracker uses an IMM Kalman Filter to project the track to the required time. In addition the track makes available the list of sensor

observations which are correlated into the track hypothesis. Although the tracker is maintaining a set of hypothetical track extensions it only publishes the extensions selected by assignment processing.

In other applications the tracker will provide a more extensive set of parameters related to the tracks. The added data available includes the complete hypothesis tree for each track, hypothesis covariance matrices and raw input data for Track Input Data.

LOCKHEED MARTIN

Tracking & Sensor Fusion Capabilities

Operator Interfaces and Output Data

- System does not generate the user display.
- Primary Output is Integrated track file and APIs to support display system's queries. Data available includes:
  - Smoothed track file Kalman Filter state variables and covariance data
  - Correlated reports with all associated data (i.e.. Track history to users specified depth)
  - Track ID, Classification, Current status
  - Alternative tracking hypotheses
- Operators can fully control tracking operation
  - Specify priority tracks
  - Control automatic or manual track initiation, Initiate a manual track
  - Specify Fault Tolerance options & Initiate FT recovery

Federal Systems
Owego, NY

Operator Controls

Certain aspects of the surveillance process are well suited for automation while other aspects require significant operator involvement. The operator controls permit the tracker to be used as a tool by the operator. For example, the tracker includes the functions necessary to initiate and terminate tracks. But operators don't always want to delegate this task to the tracker, or perhaps they want full control in one region while delegating the control for other regions. As is illustrated by the list of operator commands in the table the SW provides this varied capability.

One significant enhancement provided is the capability for the operator to control how processing is allocated to tracks. Sensor reports will be displayed to the operator within three seconds of when they are received. This delay is split equally between signal processing, tracking and display processing. To meet the need for tracking within one second of data receipt from signal processing the tracker segments the surveillance domain and uses limited initial processing for cluttered regions. The tracker then catches up with normal processing by utilizing the excess processing capability associated with latter sparsely populated surveillance segments. Operator controls permit management of this process to best satisfy mission needs. The size of surveillance regions can be specified and specific tracks can be given high priority. High priority tracks will be processed first so this capability serves to focus the initial processing on results that are of most interest to the operator.

LOCKHEED MARTIN

Tracking & Sensor Fusion - System Features

No. and Title	Description
1) High Performance K-dimensional Assignment Solver	Permits optimal solution based upon sensor kinematic and characteristic data plus database information
2) Segmented Sector Processing	Permits operator output in minimal time
3) Realtime Load Controls	Tracker converges upon optimal solution. Controls permit definition of deadlines and catch up periods.
4) Static Memory Allocation	Optimize runtime performance by compiling for maximum system resources
5) Dynamic Hypothesis Tree Size	Manages the size of the hypothesis tree to fit available processing and memory resources
6) Track Priority Queues	Organizes tracks to allow for Realtime Load Controls
7) High Priority Tracking	Specifies tracks that may use more processing resource if needed.
8) Tracker Multi-threading	Permits asynchronous IO processing.
9) Origin Management	Adjusts for origin of sensors
10) Specialized Sensor Data Processing	For current sensors we include sensor specific processing (ex. IFF code variation rules) and attachment points for future sensor specific processing modules.

Federal Systems
Owego, NY

LOCKHEED MARTIN

Tracking & Sensor Fusion - System Features

11) Fault Tolerance Extensions	Checkpoint and data recovery support
12) Open System Interface Support	Allows for alternative IDL data descriptions in CORBA systems or support for non-CORBA systems via replaceable IO module
13) Training Support	Simultaneous support for live, simulated and exercise tracks
14) Operator Controls	Support for: a) Initiate, Split, Merge, Drop or Prioritize Track b) Define automatic track initiation regions c) Enable specialized sensor data processing (i.e. multipath) d) Change origin e) Set operational mode – ex. Enable Training Modes f) Upload tracks for checkpoint recovery
15) Integrated Target Classification	Classification hierarchy based system that will, for each track hypothesis, map the sensor characteristic data to classification tree and based upon the mapping designate hypothesis classification based upon most resolved and supported category.

Federal Systems
Owego, NY

A number of features relate to how tracking is managed in the real time environment. In each case the tracker SW has been developed to provide what is believed to be suitable processing, but it also provides the means for operators to override that processing. For example the tracker will normally adjust the size of the hypotheses tree extension window if faced with overrun situations. Operator functions permit the user to make this decision explicitly.

Testing

Program testing has occurred in several distinct phases. The SW was originally developed as a LM IRAD project. To support development, tools were created to model the environment and the sensors. Throughout the development process those tools were used to generate testcases and a large set of regression tests. To automate testing, measures of effectiveness (MOE) were developed and automated.

When the Air Force executed the Best of Breed Tracker selection, they also developed a large set of testcases and MOEs. These were made available to developers after the LM tracker was selected the winner. LM developers used these testcases to further refine the test suite available.

As a result of the BBT win, LM was awarded a contract to recode the algorithm and develop a production product. In this phase three distinct testing stages were identified. As a result of the extensive testing done by LM prior to delivery of the product we have found few program issues requiring fixes throughout the qualification and flight test period.

LOCKHEED MARTIN

Tracking & Sensor Fusion Capabilities

Production Testing

Three levels of testing were utilized

- LM Formal Qual (IVV) - for this testing we developed full environment and sensor modeling tools and a full set of evaluation metrics and displays
  - full set of testcases were developed by an independent tester
  - results were continually feed back into development process
- System Qual - it was started 5/98 and is continuing
  - to date this has resulted in few if any issues
- Flight Testing - started in 7/98 and continuing
  - primary issues have related to how application/display SW interacts with tracker
  - operators are requesting more detail than what application is requesting from tracker

Federal Systems
Owego, NY

Future Plans

The tracking and sensor fusion processing developed was not developed specifically for AWACS. We focused on a rigorous and theoretically based algorithm development. This has been beneficial in several ways. Because the process was not tuned to specific testcases we were not impacted adversely when the nature of the target environment changed. With a general approach to sensors we can extend the set of supported sensors without completely redoing the processing. For example we are currently working to add support for sonar observations. Finally when we were asked to extend the tracker to provide track id we were able to take an existing id/classification algorithm developed by another site and integrate it with the existing tracker. Further this integration was accomplished in the space of only a few weeks. The following section describes how this can be done and explores some of the potential extensions.

LOCKHEED MARTIN

Tracking & Sensor Fusion Capabilities

Future Plans & Pricing

- Current work is focusing on integrating MSI, Sonar, and off-board track data
- Assignment Solver allows factors other than just kinematic track fit to be evaluated.
  - We will use that characteristic to extend tracking to ID and Sensor Resource Scheduling/Planning
- With complete air picture we can provide effective Situation Assessment

- Limited capability versions of tracker are available for 120 day test and evaluation (Full capability versions are classified, but available)
- In low quantity we offer object versions for $40K
- Source code licensing is available
- Quantity pricing is negotiable & generally would include engineering support

Federal Systems
Owego, NY

Open System Tracking

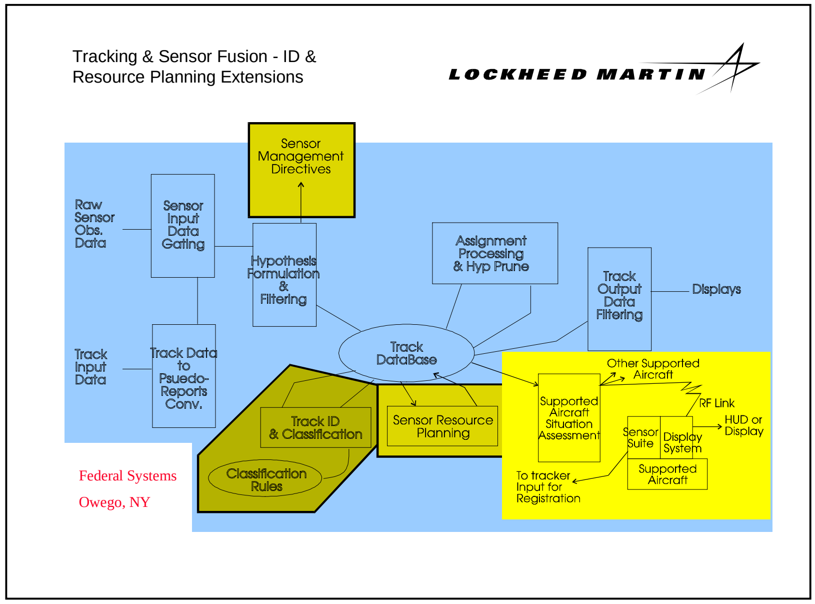

The AWACS platform is an open system distributed processing system based upon the OMG CORBA standard. The tracker fits into this environment by providing IDL interfaces for all incoming, outgoing and operator control messages.

They system is effective and it has enabled very efficient system integration. Based upon this experience we are extending this concept to provide IDL based access to the tracker's Track Database. That extension will permit other applications to be tightly integrated with the tracker, even if the tracker team does not develop those applications.

The first example of this capability is the Track ID & Classification Processing. We have already tested this approach by selecting a product developed for the Army RPA program, integrating it with the tracker, and delivering the integrated product to the prime contractor for a competitive MSI evaluation. In that evaluation the combined product satisfied all requirements in a timely manner while competitive products were unable to satisfy the time requirements.

In that evaluation we attached the ID processing as an additional output function. However, testing showed that it would have been more effective if attached directly to the database. As a result we are currently extending the database to provide the open interfaces. With the open interface we anticipate that at least two other types of situation awareness processes will be added.

In the first case we will use existing programs to analyze the tracking quality relative to the sensor performance. By establishing policies that permit the tracker to request additional data from the sensors we anticipate that a significant improvement in tracking quality during target maneuvers.

The second case appears to be more significant. If AWACS is supporting fighters it maintains a tracking history and state which is effectively a 'god's eye view' of the battle. At the same time fighters are using their own field of view radar to search and track targets. The fighter's view is limited to a pie shaped sector out ahead of the platform. It does not see what is off to the side or behind. It is possible for the AWACS platform to acquire the observations being made by the fighter and one can imagin a mechanism to return AWACS track images to the fighter. However, normal sensor errors and sampling rate variations make it unlikely that the AWACS tracks would directly match the fighter's tracks or visa versa. The fighter is tracking continuously while AWACS is sampling so the fighter sees maneuvers immediately while AWACS generally requires additional sampling intervals before maneuver hypotheses are selected. Given this background what we are proposing is to use the capability of the tracker to map AWACS tracks, or supported hypotheses directly to the perspective to the fighter and return to the fighter a properly registered image. In effect this would provide the fighter with side and rear looking radar. When this capability is combined with effective track id processing then the image returned to the fighter can highlight both targets and threats. Not only does this enhance the information available to the fighter pilot but it adds evidence to the AWACS database there by enabling AWACS operators to see maneuvers earlier.

Summary

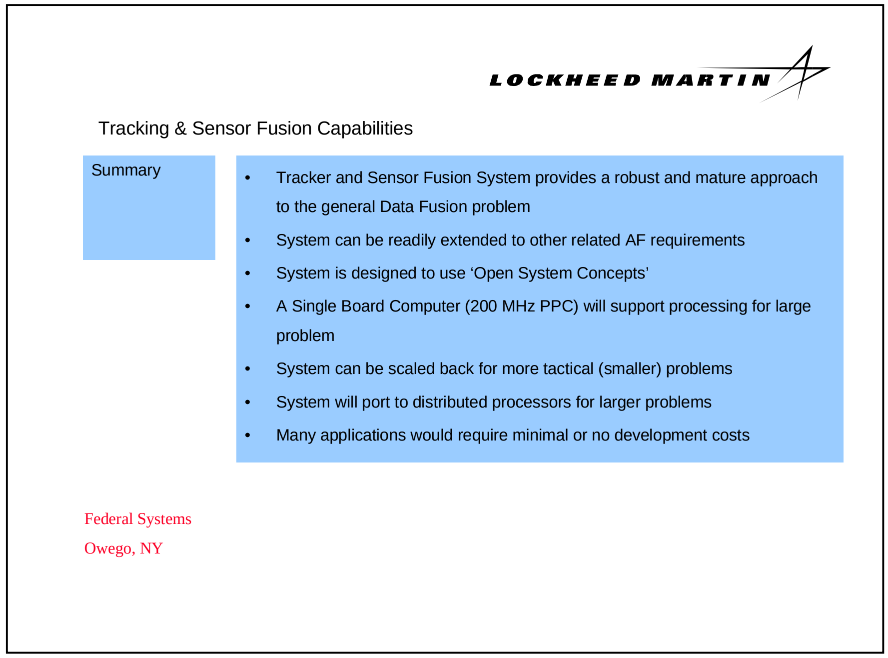

The AWACS tracking and sensor fusion software represents a significant forward step in tracking technology. It was developed to take full advantage of the growth in processing capability that has occurred in the 90s.

The existing SW is well tuned to the needs of real-time processing. Results are delivered to the display system within 1 second of when inputs are received and this performance is maintained under very heavy tracking loads.

The existing AWACS uses 200 Mhz PPC with 128M of storage for tracking. With that process we are able to track in excess of 1000 targets.

The bulk of the processing is used to perform Kalman Filtering on each developed hypothesis. This processing constitutes millions of totally independent tasks so it is completely amenable to parallel processing. Work is currently being done by Mitre Corp, Bedford MA to exploit this capability. To date there has been no need to consider the use of parallel processing for any tracking applications that have been sized for this product.